

Preparing to Run on Electra Skylake Nodes

To help you prepare for running jobs on Electra's Skylake compute nodes, this short user guide includes information on the general configuration of Skylake nodes, compiling your code, running PBS jobs, and checking allocation usage.

Overview of Electra Skylake Nodes

Electra includes 2,304 Skylake nodes, which are partitioned into sixteen physical racks. Each node contains two 20-core Xeon Gold 6148 sockets (2.4 GHz) and 192 GB of memory. The nodes are connected to the Electra InfiniBand network (ib0 and ib1) via four-lane Enhanced Data Rate (4X EDR) devices and switches for internode communication. The ib1 fabric is used mainly for I/O and is connected to the Pleiades Lustre filesystems. In addition, Electra and Pleiades share the same home filesystems, Pleiades front-end systems (PFEs), and PBS server. You can access Electra only through PBS jobs submitted from the PFEs.

Usage of Electra Skylake nodes is charged to your project's allocation on Pleiades at the rate of 6.36 Standard Billing Units (SBUs).

Compiling Your Code For Skylake Nodes

The Skylake processors include the Advanced Vector Extensions 512 (AVX-512). Intel AVX-512 optimizations are included in Intel compiler version 16.0 and later versions. We recommend that you test the latest Intel compiler, **2020.4.304**, which may include better optimizations for AVX-512.

For Skylake-specific optimizations, use the compiler option **-xCORE-AVX512**.

If you want a single executable that will run on any of the Electra or Pleiades processor types, with suitable optimization to be determined at runtime, you can compile your application using the option:

-O3 -xCORE-AVX512, CORE-AVX2 -xAVX.

Note: The use of either of these options, **-xCORE-AVX512** or **-xCORE-AVX512**, could either improve or degrade performance of your code. Be sure to check performance with and without these flags before using them for production runs.

Running PBS Jobs on Electra Skylake Nodes

To request Electra Skylake nodes, use **:model=sky_ele** in your PBS script, as shown in the example below.

Important: Because MPT 2.15 and earlier versions do not support the ConnectX-5 host channel adapters (HCAs), the environment variables **MPI_IB_XRC** and **MPI_XPMEM_ENABLED** have been disabled for jobs running on Skylake. If your MPI applications perform significant MPI collective operations and rely on having these two variables enabled to get good performance, use MPT 2.16 or newer versions. To use the NAS-recommended MPT version, use the command:

```
module load mpi-hpe/mpt
```

Sample PBS Script For Electra Skylake Nodes

```
#PBS -l select=10:ncpus=40:mpiprocs=40:model=sky_ele
#PBS -l walltime=8:00:00
#PBS -q normal

module load mpi-hpe/mpt
module load comp-intel/2020.4.304

cd $PBS_O_WORKDIR

mpiexec -np 400 ./a.out
```

Checking Allocation Usage For Electra Skylake Jobs

To track allocation usage for jobs running on Electra Skylake nodes, run:

```
acct_query -c electra_S
```

Examples:

- To track the total SBUs usage for one of your GIDs (for example, s0001) since 9/30/17:

```
% acct_query -c electra_S -ps0001 -b 9/30/17
```

- To track the SBUs usage for each of your jobs run today:

```
% acct_query -c electra_S -olow
```

For more information, see [Skylake Processors](#).

Article ID: 551

Last updated: 13 May, 2021

Revision: 27

Systems Reference -> Electra -> Preparing to Run on Electra Skylake Nodes

<https://www.nas.nasa.gov/hecc/support/kb/entry/551/>